CLAIMS

1. A system adapted to distribute route selection in an implementation of a routing protocol executing on a router of a computer network, the system comprising:

a first process of the routing protocol configured to receive announced paths from peers of the router and perform a first stage of route selection to select partial best paths;

a second process of the routing protocol configured to perform a second stage of route selection to select best paths in response to the partial best paths forwarded by the first process, the second process further configured to send the selected best paths to the first process for announcement to the peers.

2. A method for distributing route selection in an implementation of a routing protocol executing on a router of a computer network, the method comprising the steps of:

receiving announced paths from peers of the router at a plurality of first processes of the routing protocol;

performing a first stage of route selection at the first processes to select partial best paths;

forwarding the partial best paths to a second process of the routing protocol;

performing a second stage of route selection at the second process to select best paths; and

sending the selected best paths to the first processes for announcement to the peers.

3. The method of Claim 2 wherein the routing protocol is a Border Gateway Protocol (BGP) and wherein route selection includes a BGP best path selection algorithm.

4. The method of Claim 3 wherein the first processes are speakers and wherein the second process is a BGP routing information base (bRIB).

5. The method of Claim 4 further comprising the steps of:

providing a plurality of first processors configured to run the speakers; and

providing a second processor configured to run the bRIB.

6. The method of Claim 4 wherein the step of performing the first stage of route selection comprises the step of splitting the announced paths for each prefix into a plurality of groups such that within each group, the BGP best path selection algorithm is a transitive relation.

7. The method of Claim 6 wherein the step of splitting comprises the step of grouping the paths according to an autonomous system (AS) from which they were received.

8. The method of Claim 7 wherein the step of performing the first stage of route selection further comprises the step of calculating a best path in each group using the BGP best path selection algorithm.

9. The method of Claim 8 wherein the step of performing the first stage of route selection further comprises the step of performing a comparison between each best path from each group.

10. The method of Claim 9 wherein the step of performing a comparison further comprises the steps of:

selecting a path with a highest degree of preference;

selecting a locally originated path over a learned path;

selecting a path with shortest AS_path; and

selecting a path with lowest origin.

11. The method of Claim 10 wherein the step of performing the first stage of route selection further comprises the step of forming a set of partial best paths forwarded to the

bRIB from any paths that have not been discarded by running the algorithm at each speaker.

12. The method of Claim 11 wherein the step of performing the second stage of route selection comprises the step of using the full BGP best path selection algorithm to select a best path per prefix from among the partial best paths received from all speakers.

13. A system adapted to distribute route selection in an implementation of a routing protocol executing on a router of a computer network, the system comprising:

a plurality of first processes of the routing protocol configured to receive announced paths from peers of the router and perform a first stage of route selection to select partial best paths;

a second process of the routing protocol configured to perform a second stage of route selection to select best paths in response to the partial best paths forwarded by the first processes, the second process further configured to send the selected best paths to the first processes for announcement to the peers.

14. The system of Claim 13 wherein the routing protocol is a distance vector routing protocol.

15. The system of Claim 13 wherein the routing protocol is a Border Gateway Protocol (BGP) and wherein route selection includes a BGP best path selection algorithm.

16. The system of Claim 15 wherein the first processes are speakers and wherein the second process is a BGP routing information base (bRIB).

17. The system of Claim 16 further comprising:

a plurality of first processors configured to run the speakers; and

a second processor configured to run the bRIB.

18. The system of Claim 17 wherein each speaker splits the announced paths for each prefix into a plurality of groups such that within each group, the BGP best path selection algorithm is a transitive relation.

19. The system of Claim 18 wherein the groups are organized according to an autonomous system (AS) from which they were received.

20. The system of Claim 19 wherein each speaker further calculates a best path in each group using the BGP best path selection algorithm.

21. The system of Claim 20 wherein each speaker further performs a comparison between each best path from each group.

22. The system of Claim 21 wherein the speaker performs the comparison by (1) discarding the path with the lower degree of preference, (2) discarding a learned path if the other path is locally originated, (3) discarding the path with longer AS_path, and (4) discarding the path with higher origin.

23. The system of Claim 22 wherein any paths that have not been discarded by running the algorithm at each speaker form a set of partial best paths that are sent to the bRIB.

24. The system of Claim 23 wherein the bRIB performs the second stage of route selection using the full best path selection algorithm to select the best path per prefix from among the partial best paths received from all speakers.

25. Apparatus adapted to distribute route selection in an implementation of a routing protocol executing on a router of a computer network, the apparatus comprising:

means for receiving announced paths from peers of the router at a first process of the routing protocol;

means for performing a first stage of route selection at the first process to select partial best paths;

means for forwarding the partial best paths to a second process of the routing protocol;

means for performing a second stage of route selection at the second process to select best paths; and

means for sending the selected best paths to the first process for announcement to the peers.

26. A computer readable medium containing executable program instructions for distributing route selection in an implementation of a routing protocol executing on a router of a computer network, the executable program instructions comprising program instructions for:

receiving announced paths from peers of the router at a plurality of first processes of the routing protocol;

performing a first stage of route selection at the first processes to select partial best paths;

forwarding the partial best paths to a second process of the routing protocol;

performing a second stage of route selection at the second process to select best paths; and

sending the selected best paths to the first processes for announcement to the peers.